

Advanced C4ISR Platforms with Machine Learning Capabilities

Fred Maymir-Ducharme, Ph.D., Executive Consulting Architect

IBM Federal, CTO Office
6710 Rockledge Drive
Bethesda, MD 20854
UNITED STATES OF AMERICA

FredMD@us.ibm.com

Richard Ernst, Executive Client Architect

IBM Federal, CTO Office
6710 Rockledge Drive
Bethesda, MD 20854
UNITED STATES OF AMERICA

remst@us.ibm.com

ABSTRACT

C4ISR environments continue to grow in size and complexity – compounded by the need for faster actionable intelligence. NATO requires leading edge platforms to ingest large volumes of data, provide low latency, advanced analytics to automate and correlate complex relationships and observables in the data. This paper & presentation will provide an overview of a platform that currently: processes 400 TB of metadata daily from widely diverse, worldwide, geospatial and meteorological content stores; analyses and persists over 100 Peta Bytes of data; responds to 200 million individual user queries, and generates over 26 billion analytics weather forecasts daily across the globe (supporting 15-26 billion API requests daily.) The authors will describe how this platform can be coupled with a Machine Learning framework, enabling Data Engineers, Data Scientists, Mission Analysts and Developers to exploit C4ISR streaming data (data in motion) to enhance the capabilities of C4ISR Mission Operators.

1.0 INTRODUCTION

The military concept of C4ISR (Command, Control, Communications, Computers, Intelligence, Surveillance and Reconnaissance) represent the key elements of Modern Defence Systems. At its simplest level, the C4ISR model illustrated in Figure 1 below depicts the Command & Control managing the military's Land, Naval, Air & Space components, exploiting Intelligence from its Information, Surveillance and Reconnaissance platforms. Implicit in the model is the reliance on Communications; and of course, as the size and complexity of military operations have increased, so has the dependence on Computers. Another observation is that as messaging technology matured (e.g., transitioned from analog to digital messaging), networks capable of transmitting audio, video and all types of data, (structured, unstructured, streaming, etc.) coupled with interoperable systems at multiple layers added to the complexity. The sensors, networks, and systems that comprise these C4ISR environments continue to grow in complexity and scale. The amount of data that these sensors are now capable of generating continues to grow. When coupled with the interoperability provided by current networks and interoperable systems, the total amount of data that is available for exploitation continue to grow at an alarming pace. At the same time, the tactics and techniques of asymmetrical warfare, and the modernization of adversaries have compounded "The Fog of War", increasing the need for both operational tempo and maximizing effects of Joint operations (land, sea, air & space).

The complexity of current C4ISR environments, the sheer volume, variety, and velocity of the generated or collected data, and increased operational tempo have given rise to the need for automated systems that can

ingest data at high volume, process it with low latency, and derive actionable intelligence from all of it, providing the ability to gain the advantage on the battlefield. This scenario is exacerbated with the proliferation of unmanned systems (e.g., UAVs).

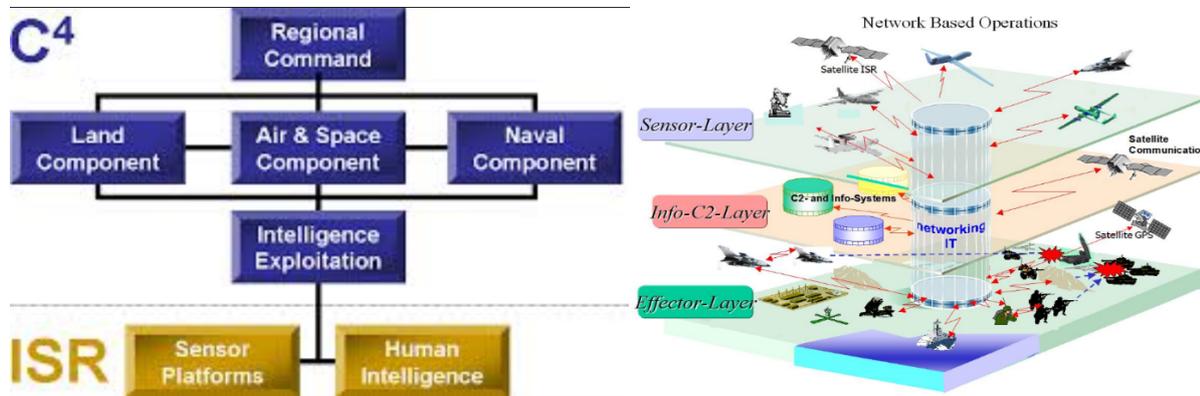


Figure 1 C4ISR Model and Network-Centric Warfare

Today’s leading-edge platforms, provide the ability to ingest large volumes of data (structured, semi-structure, unstructured), across a spectrum of streaming data, data at rest, and batch processing. These systems must provide in-memory, low-latency processing and advanced analytics capabilities to automate and correlate complex relationships and observables across all of the data being provided by the myriad of purpose-built systems. [1]

While the analytics industry has made great strides in analytics technology (e.g., Descriptive, Prescriptive, Predictive), the quest for Cognitive Analytics persists. Cognitive Analytics has often been associated with Artificial Intelligence (AI), which entails the automation key human intelligence tenets: Understanding, Learning, and Cognition. [2] Besides the advances in analytics, the explosion of data and information produced and available today (Big Data) has been a key enabler for Cognitive Computing. Figure 2 below illustrates the inclusive relationships between AI, Machine Learning (ML) and Deep Learning (DL).

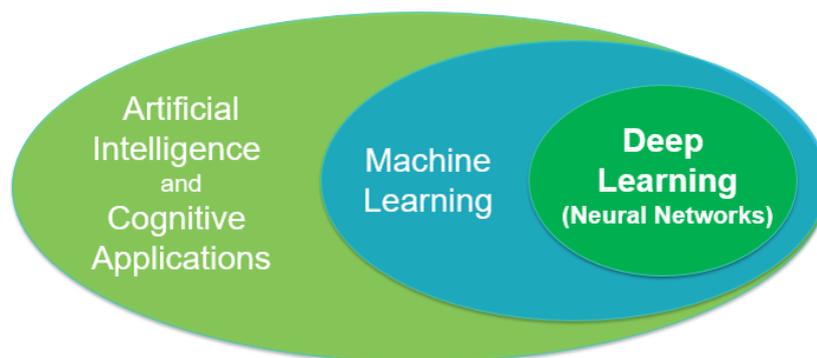


Figure 2 Artificial Intelligence and Machine Learning Perspectives

ML goes beyond traditional analytics, providing the ability to dynamically modify and extend the analytics’ original capabilities to improve results and desired outcomes as new data is ingested and processed. ML is an evolving technology. There are numerous approaches to designing, architecting, implementing, and optimizing ML systems. Given the very large data volumes and empirical data available to NATO organizations, ML represents a very attractive approach to optimizing and extending existing capabilities and developing new ones (e.g., early threat anticipation and resolution, extended situational awareness.)

2.0 High Volume Low Latency Platforms

The proliferation of sensors and sensor platforms across the land, naval, air and space components has created challenges with respect to data ingest, fusion, access, and the automation of correlating objects and activities resulting in accurate and effective activity-based intelligence.

Data Ingest: Modern C4ISR platforms need to be able to ingest and distribute a broad spectrum of types and sources of data. Data types range from structured, to semi-structured, to unstructured data. The variety of systems and sources generating and or providing access to data requires a data platform that supports standards-based interfaces and API's, as well as have the ability to connect to proprietary interfaces and protocols. We've briefly discussed variety of data. Data velocity is another dimension of data ingest, and these modern platforms must possess the ability to connect and ingest data across the entire spectrum of velocity spanning streaming data, periodic or batch, to pulls or pushes on demand. The architecture of modern platforms needs to support the ability to add or remove a data source without having to redesign or redeploy the platform itself. The use of API's and micro-services in an open, modular, and loosely coupled architecture are imperative to fielding systems that are cost effective and responsive to changing mission, and hosting environment needs.

Data Management: Stewardship and or custodianship of data (quality, policy, access, storage, integrity) must, in a real world, operational sense, be addressed by any platform designed and implemented that will ingest and produce derivative data. The platform should have provisions for creating and managing meta-data, for data access services that can broker access to specific data elements based on user or system credentials, as well as the ability to build and maintain a data model that can be accessed by data conditioning, data access, data export and data storage services.

Data Exploitation: Data is the resource that is consumed by, and fuels analytical models, in turn, yielding information, knowledge and insight. C4ISR platforms of the future will need to provide access to data processing in high volume, low latency run time environments. Low latency processing of data is vital to systems that support time sensitive applications and or services. The value of object data found in modern C4ISR systems are time sensitive. Examples include air or ground tracks of a threat or target.

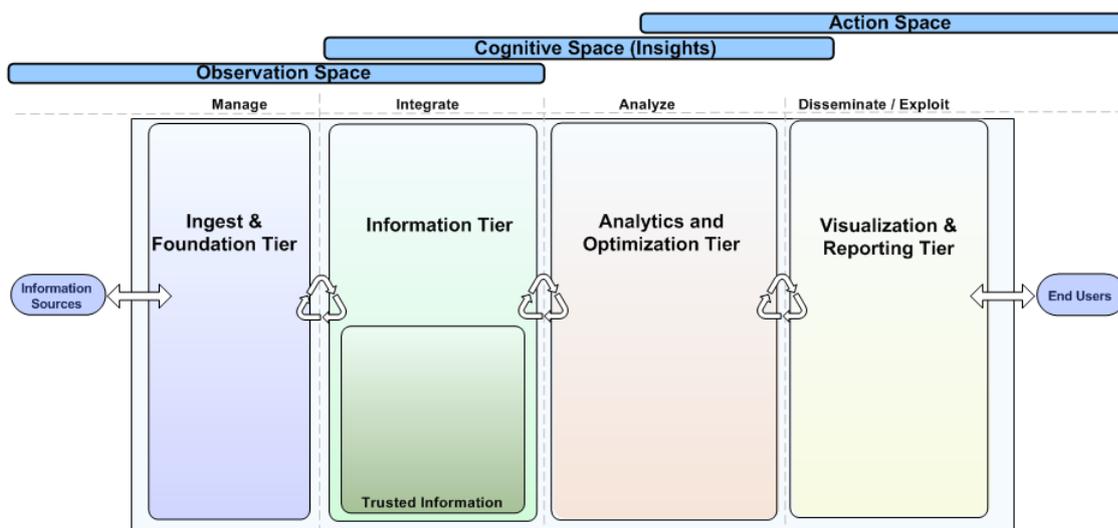


Figure 3 Observation, Cognitive, and Action Space Overview

Figure 3 above, provides a high-level illustration outlining how a modern, high volume, low latency data and analytics platform attributes correlate to the C4ISR Observation Space, Cognitive Space and Action Space (C4ISR Effector Layer). Threat or target track data must be correlated across sensors and processed with enough agility to be able to make informed decisions and then take action in order to realize an effect.

If data is the resource that fuels analytical models, then modern data processing platforms need to possess attributes conducive to data ingest and preparation. Understanding the business or mission problem to be solved, and then ingesting, transforming, cleansing, normalizing, and conditioning data can typically consume up to 80% of an effort from a time perspective on an analytics project. The other 20% is focused on developing, testing, and scoring the analytical model. If Machine Learning (ML) is required, then balance between data preparation and building the analytical model is more equitable, as the establishment of ML training environments such as Tensorflow, can very time consuming, in addition to the time required to actually perform the supervised and or unsupervised machine learning itself.

An additional factor to consider is that there are several roles required in the performance of Data Science and Machine Learning. The key roles are:

Data Steward: role that incorporates processes, policies, guidelines and responsibilities for administering organizations' entire data in compliance with policy and/or regulatory obligations. A data steward may share some responsibilities with a data custodian.

Data Engineer: data professionals prepare the “big data” infrastructure to be analysed by Data Scientists. They are software engineers who design, build, integrate data from various resources, and manage big data. Then, they write complex queries on that, make sure it is easily accessible, works smoothly, and their goal is optimizing the performance of their company’s big data ecosystem. [3]

Data Scientist: Data Scientists apply statistics, machine learning and analytic approaches to solve critical business problems. Their primary function is to help organizations turn their volumes of big data into valuable and actionable insights. [3]

Business/Mission Data Analyst: someone who analyses an organization or business domain (real or hypothetical) and documents its business or processes or systems, assessing the business model or its integration with technology. This role merges C4ISR mission knowledge and expertise with the integration of technology and analytical models to ensure that the goals and objectives are correct and properly evaluated.

Application/Service Developer: application and service developers create (program), integrates and tests applications and micro-services software. Each role has a specific skill and responsibility. Additionally, these roles need to be able to work together efficiently.

Modern Data Science Platforms provide tools that these roles are used to working with as well as the ability to collaborate and share data sets and analytical models, perform machine learning activities on supporting infrastructure, and to version and configuration manage the work products.

2.1 The Weather Company Platform

The IBM Fast Data for Business platform (FD4B) evolved from The Weather Company’s (TWC) open and modular platform. It is an example of a high-performance platform to ingest, condition, distribute, persist, archive, analyse, process, exploit, and disseminate boundless amounts of data and information. The first generation FD4B platform was developed using open source technology and runs on Amazon Web Services Cloud (AWS) with 100% system availability since going live in 2013. It provides commercial geospatial

and location-based services, and Data Access Profile (DAP) solutions at a comparable scale to Google, Facebook, and Twitter. FD4B uses data to provide analytic services that help organizations understand their data to make more informed decisions, i.e., relationships between target observables and actor(s) behaviour.

Only Google, Facebook, and Twitter have comparable scalability and performance. However, those platforms are restricted to internal business use only, whereas FD4B is proven, and available for use in ingesting, and exploiting your data for your mission needs. Figure 4 below provides an overview of the platform. The platform is designed to handle high volumes of data, process with low latency and maintain high availability. It is modular and open.

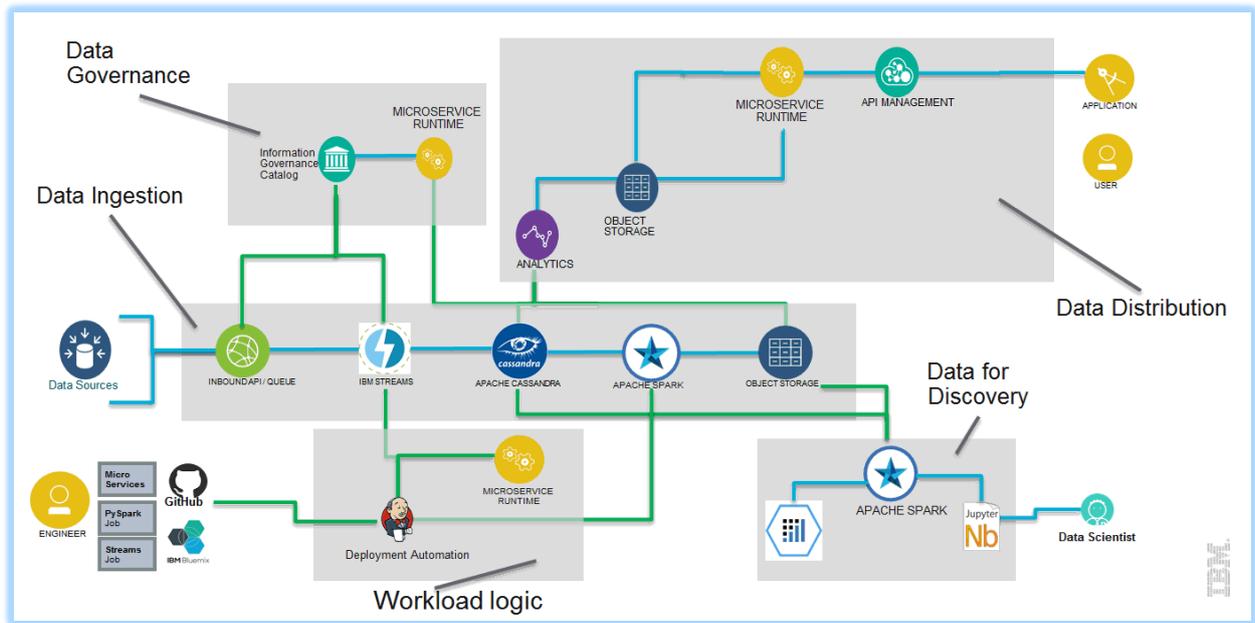


Figure 4 - Fast Data for Business Platform Overview

Platform Scalability: Using our TWC example, every day the FD4B platform delivers over 15 billion forecasts, processes more than 400TB of metadata from widely diverse, worldwide geospatial, and meteorological content stores, responds to 200 million individual user queries in the cloud, and delivers content to a diverse array of user applications and platforms resulting in 15-26 billion daily API requests. The strategy is to be within 100 milliseconds of everyone on the planet and leverage the global reach and availability of Cloud computing to deliver on that strategy. This same approach can be applied to C4ISR environments in the air, or on the ground. The approach scales up or down as needed.

Data Scalability: Weather is a volatile dataset that changes every 15 minutes or less. To scale quickly on a global basis, TCW pushes over four gigabytes per second of forecasting data through AWS services. For forecasting, the FD4B platform handles massive data movement, replication, and data distribution in an AWS cloud-based, distributed environment. Historical data from satellites, radars, forecast models, users and weather stations is stored and persisted.

Extensibility: The platform enables clients to expand the framework with new API and modules that can be integrated into workflows. This is an important point because the C4ISR mission and technology are always evolving, i.e., new algorithms to process SIGINT, ELINT, etc. New analytics APIs can be added to perform analytics on the data as its ingested to streamline decision making. For its weather business, TWC runs 150 analytics models in parallel. Using these automated analytics, it can run its global business with less than 150 analysts. C4ISR specific models can be developed, and then deployed into operational run time environments.

Data Sources: Today, the platform ingests information from more than 150 different sources. New data type can be added in a few days. Data is stored leveraging the efficiency of cloud storage as opposed to an on-premises storage solution and eliminating the hassle of managing a storage platform. The platform natively supports a polyglot storage model, so other storage mediums are acceptable and can be easily added.

API Driven Platform: FD4B is an API driven platform so that clients can build applications tailored to their mission and operating environment. FD4B developers worldwide take advantage of the API-driven platform to build applications and workflows for their customer’s specific use cases. Any developer can use the API to develop services for the data platform, other cloud environments, or on-premises hardware. The platform is robust enough to handle between 15 and 26 billion transactions each day at 173,000 to 300,000 transactions per second, depending on the weather which drives user demand. The weather API is one the world’s most-used public API’s.

2.2 Data Science Platform

Interest in, and adoption of data science and machine learning is booming with a profusion of new tools, new algorithms, new libraries, and new data sets such as Weather, Twitter and more. Data-driven companies such as Uber and Airbnb have shown the rest of the world the impact world-class data science can have on business results. The challenge companies have to extract business insights from their data is exacerbated by a continued shortage of skilled data scientists. That skills shortage, coupled with the fact that many existing tools are siloed, requiring different skill sets, has made it increasingly difficult for companies to derive valuable insights that could drive decision-making and business results.

The Cross-industry standard process for data mining (CRISP-DM) still provides a good model as we think about understanding the business/mission, the data that supports it, and how to condition it, create analytics and visualizations that support extracting value from it (see figure 5).

An example new data and analytics platform that is focused on Data Science is Data Science Experience (DSX). DSX is a platform designed from the ground up for the Data Engineer, Data Scientist, Mission/Business Analyst, and Application Developer.

DSX provides end-to-end management of the machine learning workflow through a suite of tools and capabilities that enables data scientists to accelerate their productivity and keeps models deployed across the enterprise current. DSX Local is available on premises and offers the following benefits:

- Integrates the most popular open source data science tools (for example, Jupyter Notebooks, Zeppelin Notebooks, Spark, RStudio, SPSS, and more) in a single location and enables access to common data and framework, eliminating the problem of siloed tools.
- Scales data scientist productivity through collaboration, enabling data scientists to tap into the expertise of their peers and to collaborate over specific multiple projects.
- Offers optimization capabilities by embedding optimization engines such as IBM CPLEX and IBM CP Optimizer with Python API as a modelling interface.
- Allows data scientists to develop models with a wide range of algorithms that can be trained and improved with machine learning services. Models can be rapidly deployed to production and continuously monitored by data scientists in a single model management dashboard.

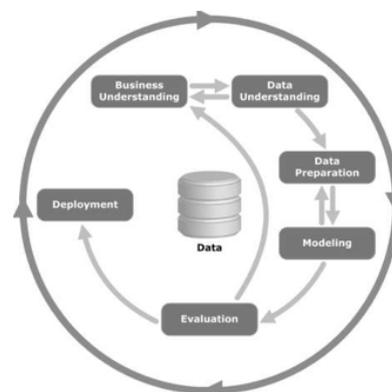


Figure 5 CRISP-DM Process

DSX is an interactive, collaborative, cloud-based environment where data scientists can use a growing set of tools to activate their insights such as RStudio, Jupyter, Python, Scala, Zeppelin, Spark, SPSS, IBM Watson Machine Learning, and more.

3.0 Machine Learning

AI has been a computer science field since the 1950s, with a focus on simulating human understanding and reasoning, and automating cognitive cybernetics. [2] AI's early approaches entailed creating sets of rules (e.g., "conditions" mapped to "function") which could be weighted, nested (recursively), subset, or aggregated to support complex multi-relational decision paths. As the size and complexity of the rules base grew, the number of permutations quickly became intractable – given computing limitations at the time. Optimization advances were enabled by modelling techniques such as Petri-Nets and Semantic Nets – which ultimately shifted their focus towards the computer science field of Analytics. Dr. Callaos clearly distinguishes between Informatica (Analytics) and Cybernetics. [4] Advances in computer processing, and the ability to manage large volumes of heterogeneous data have revitalized the field of AI, particularly with the emergence of ML, which leverages both, informatics (analytics) and cybernetics (e.g., decision support or autonomous computing.)

As described in the Introduction, ML includes a variety of analytic components, adding algorithms that "learn" and improve desired outcomes. ML algorithms essentially provide data analysis approaches that leverages additional information or extracts additional knowledge beyond those of the underlying analytic components and data models. One of the simplest examples of ML is an automatic categorization system. The traditional categorization (analytics) tool may be a rules-based system that categorizes documents based on the inclusion, exclusion, and cardinality of keywords. The ML version may then allow a human (or external system) interface that changes the document binning – the results of which are then used to modify the previous rules to better reflect the modified results.

Yes, ML has been around for many years, but the explosion of data and advances in analytics that can be applied to a plethora of data types (structured and unstructured text, multimedia, and non-traditional data types such as sonar or radar) has made for many advances in the field. At a high level, ML systems are composed of three key platforms: one to input the data, one to analyse and learn, and the third to present or act on what was learned. The first and third component can be combined (e.g., input data elements and data output elements) or may be separate (e.g., in the case where results are used for autonomous or semi-autonomous action). [5] The Data Platform described in Section 2 is an example of a platform that combines the first and third platforms -- which support the acquisition, ingest, management and storage of information, as well as elements to publish or disseminate information. Today's advanced data platforms allow us to exploit a plethora of different data types, and at a much grander scale than was previously possible. This section (Section 3) provides an overview of Machine Learning platforms, which include components for analyses and components for "learning."

Before a ML system can "learn," it must be able to analyse... The analyses provide methods to evaluate or measure different types of data (e.g., measure the speed of a car), and to provide one of more assessments – which will later be improved by learning techniques that validate or invalidate elements of the analyses. For example, the speed of a car may be analysed as "speeding" if over 25MPH in a neighbourhood; but the same speed (say 30MPH) may not be "speeding" if on a highway. Contextual analytics can be used in the latter example to "regulate" safety. But as one learns later in life, there are many other aspects to neighbourhood or highway auto safety than the velocity. Children playing outside in the neighbourhood mid-day, or heavy volumes of commuters entering or exiting the highway can provide just as important perspectives when it comes to evaluating speeding. Section 3.1 describes an exemplar taxonomy of analytics that are used in many of today's ML systems.

There are two primary categories of learning techniques (Supervised and Unsupervised) – which can then be further broken down into sub-categories. Just as a note, there are other categorization schemes, which

raise some sub-categories (e.g., Reinforced) into primary categories; but the authors prefer the simpler two category taxonomy for discussion purposes. Supervised techniques include some sort of initial or pre-determined guidance (e.g., an initial categorization scheme, or information labelled with associated outcomes deemed correct); whereas Unsupervised techniques have no pre-conceived models or meta-data, and expect the learning to come from the analysis of the data.

For example, using the categorization example described earlier, the Supervised approach would generally start with an initial set of rules or a preliminary model for categorizing the documents; whereas the Unsupervised approach might start without either, and apply learning techniques to discover groupings (e.g., events that co-occur spatially or temporally) and categorization schemes (e.g., variables and ranges that result in relatively distinct groupings.) These are described in Section 3.2.

As Figure 2 illustrates, there is a unique subset of Machine Learning that is called “Deep Learning.” Deep learning is typically depicted by the use of high performance platforms to improve the speed and accuracy of the machine learning techniques. This is accomplished by exploiting very large volumes of data as inputs into highly parallel systems. These systems use General Processing Units (GPUs), General Purpose GPUs (GPGPUs) and Field Programmable Gate Arrays (FPGA) with very fast ingest and I/O mechanisms to process more data in a short amount of time than is possible with traditional, sequential platforms (e.g., Von Neumann architectures). These can be broken down into Single Instruction, Multiple Data (SIMD) architectures such as Pipe & Filter (vector technology) architectures that can “scale out,” or Multiple Instruction, Multiple Data (MIMD) architectures such as the NOA’s super-computers that can also “scale up” with large shared memory capabilities.

Deep learning is very applicable to imagery (e.g., image recognition and various modes of biometric systems) as well as to non-linear heuristics such as encryption & decryption. Space and time constraints limit our covering deep learning in detail.

3.1 Analytics

The field of analytics is rather broad; and there are many taxonomies and classification schemes to describe the wide range of analytic techniques and approaches. The following is one such exemplar categorization of analytics, which incrementally build on each other for advanced capabilities.

- **Descriptive analytics** deal with data from the historical perspectives. E.g., How does my data map to historical models and metrics? This includes financial management systems, regulatory and compliance systems, etc.
- **Predictive analytics** deal with the future perspectives. E.g., What can we anticipate will happen, given historical trends (temporal, co-occurrence & relational). This includes statistical systems, weather systems, etc.
- **Prescriptive analytics** deal with perspectives focussed on specific outcomes. E.g., how can we maximize profits, or reduce the time to market. These include contingency planning systems, decision support systems, etc.
- **Cognitive analytics** deal with multiple perspectives. Cognition includes understanding, reasoning, and action. E.g., combining contextual models (data within relevant context) to understand and reason about a situation, augmented by prescriptive analytics that take multiple perspectives into consideration before acting or recommending an action.

3.2 Machine Learning Techniques and Platform

ML combines analytics (such as those described in the previous section), as well as various techniques. Orthogonal to the categorization scheme in Section 3.1, there are analytic techniques (e.g., statistical, classification, clustering, etc.) that are applied in ML systems to add understanding and reasoning capabilities, and in addition to the analytics, ML uses a variety of “learning” techniques.

Supervised learning includes a variety of techniques, these include:

- **Semi-supervised learning** – where the “training data” is partially labelled or mapped to incomplete models. The learning algorithms identify (or discover) additional labels (metadata) that help refine results and may extend the model(s) models based on newly learned attributes beyond the base model variables. E.g., A classification scheme that incrementally improves the soundness and completeness of the models by varying the ranges in the model variables, or by adding new variables that improve the fidelity of the results.
- **Active learning** – where the intended labelling or specific data for analysis is constrained for the learning, in order to meet very specific and managed results. Depending on the level of interaction, the user may interact with the learning to consider adding labels or data types to the active learning. E.g., a Neural net used to predict financial results, where the number of stocks (or amount of stock company information: nodes) or complexity of the weighting (arc strength) is calibrated incrementally through longer time periods.
- **Reinforced learning** – where only a subset of the labels or models are initially applied. The user (or automated reinforcement algorithm) incrementally adds labels or extends the models (positive reinforcement for success,) or reduces the labels and complexity of the models (negative reinforcement for failures) to achieve the desired outcomes. This approach is useful when there is a major focus on the number of “false positives” or “false negatives” in the results. E.g., a Decision tree used to optimize evaluating potential medical illness.

Unsupervised learning includes various techniques, including:

- **Affinity discovery learning** – where the data is grouped in as many “relevant” ways as the processing is capable of, to determine potential groupings leading to the intended outcomes.
- **Pattern discovery learning** – where a sequence of seemingly random events or measurements can be mapped to valuable outcomes. E.g., A merchant sales maximizing system that discovers that males buying diapers and other baby necessities will often also pick up “fast food & drink” (e.g., chips, nuts & beer); hence, placing these extra purchase items between the baby aisle and the check-out counters maximizes sales.
- **Association discovery learning** – where the system stochastically discovers new aggregations and associations, without a pre-determined categorization or modelling scheme. E.g., the system may take a data set and applying a variety of affinity analysis heuristics, discover new clustering’s and aggregations.

It should be noted that as in most attempts to define a sound and complete taxonomy, there are always exceptions and hybridizations. Hence, there are various ML systems that combine supervised and unsupervised techniques, as well as exploit multiple analytics to provide additional insights (e.g., mash-ups and knowledge graph technology). Hence, the platform selected must provide a high degree of interoperability, and scale to meet the exponential data growth industry is experiencing.

Figure 6 below represents an exemplar ML platform, which provides for a hybrid, multi-cloud environment to deal with technology advances such as GPUs, FPGAs, and other emerging infrastructure technology. The platform should also support the Information Management (data lifecycle) with tooling for: data ingest; information transformation and enrichment; object, block, and file system storage, etc. The exemplar system below supports many of the AI components discussed in this paper – upon which various industry-specific solutions and AI eco-systems (e.g., DSX & Watson Machine Learning) are built.

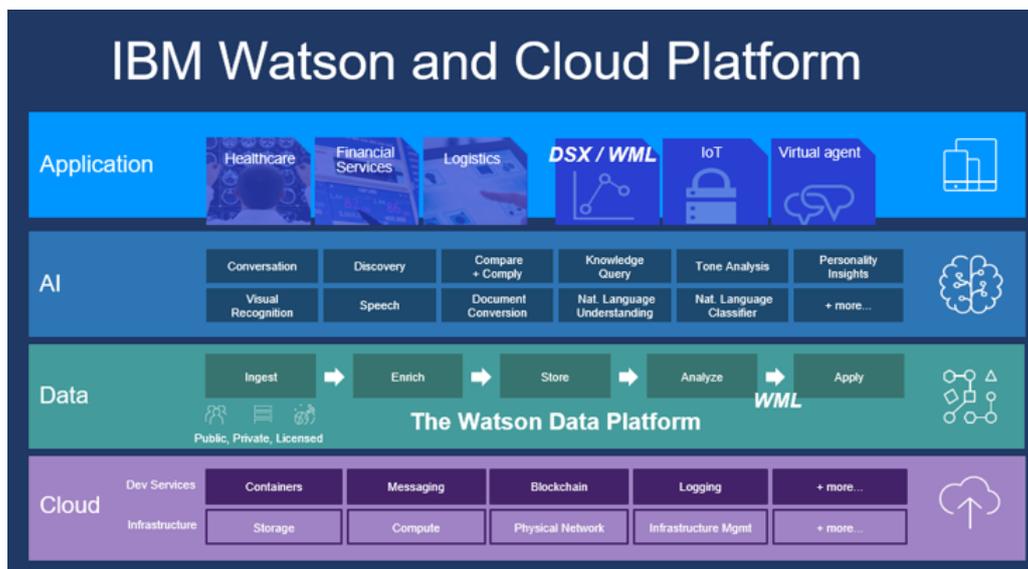


Figure 6 IBM Watson ML Platform

4.0 Summary and Conclusions

AI approaches combine Data Science, Analytics, and Machine Learning technology, which provides new capabilities across C4ISR domains. Current advances in analytics, alongside the convergence of AI and ML techniques will continue to underpin high value mission automation and decision support systems. These cognitive systems can combine natural language processing (NLP, e.g., speech-to-text, natural language queries) translangual technology (machine translation, multi-language analytics), hypothesis generation & evaluation, and continuously learning capabilities for faster, more powerful and intelligent defence systems.

NATO’s adoption of these AI capabilities requires focusing on five key areas:

1. Adapting AI technology to C4ISR and other NATO Mission Areas. An AI program can use NLP, cognitive analytics, and ML validation techniques to better understand, reason, and provide better decision support or autonomous actions. AI technology can help transform the interactions between NATO C4ISR analysts and NATO Defence operations, bringing together the relevant strategic and tactical perspectives that humans naturally try to take into consideration in their complex decision-making processes.
2. Applying cognitive analytics to computational models that reflect the international mind-sets of NATO constituents. NATO leadership and NATO members have both, common and variant perspectives – e.g., defence, political, social, and economic frames of reference that must be considered in the decision-making process. The reasoning capabilities can be used to learn from previous NATO engagements and provide member countries without the history or empirical data for relevant perspectives to consider when facing unplanned or unanticipated challenges.

3. Converging AI and NATO mission analytics. Traditional C4ISR mission analytics do a good job with “known knowns” (e.g., contingency plans for previously dealt with challenges) and with “known unknowns” (e.g., contingency planning for envisioned challenges related to previously dealt with challenges.” But the real challenge comes with “unknown unknowns” – when new and previously unimaginable challenges arise. Human cognition uses many sensors (eyes, ears, nose, fingers, etc.) as inputs for understanding, reasoning, and eventually taking action. ML provides a novel approach to processing data from non-traditional sensors for which automated support does not exist, or lags behind the traditional data processing capabilities in heritage and legacy NATO systems. Deep computing techniques can be applied to a variety of non-traditional and proprietary data sources.
4. Incorporating Analytical and Operational Intelligence into existing and evolving C4ISR systems. Today’s C4ISR systems are complex, often requiring a diverse set of software and hardware components with varying levels of integration or interoperability. Developing, designing, implementing, monitoring and optimizing the numerous elements of C4ISR systems is a grand challenge. ML has been applied to extend many of the key tenets of Enterprise Systems Management (e.g., managing servers, storage, applications, networks, security, performance, etc.) to broader systems (e.g., Intelligent Buildings or Intelligent Cities). ML is an enabling technology that can be used to exploit the Internet of Things (IoT) and other information sources that could provide NATO essential and invaluable perspectives and validation sources.
5. ML is fuelled by data. Tomorrows C4ISR platforms will need to ingest vast amounts of data (structured and unstructured), govern it appropriately (policy, access, management), and provide in memory, low-latency processing capabilities. Additionally, these High Volume, Low-Latency platforms will need to lower the barriers of entry for Data Science, Machine Learning and Artificial Intelligence.

There are many examples across Academia, Commercial Industries, and the Government in which Machine Learning is being applied at various levels of automation and autonomy. The ML Cybernetics figure below (Figure 7) illustrates additional, traditional adoption model for ML. NATO’s C4ISR systems have reached a high level of descriptive, predictive and even prescriptive capabilities. The authors envision ML will enable extending these capabilities to exploit potential advances in Intelligent Decision Support Systems, and ultimately in Intelligent Autonomous Systems.

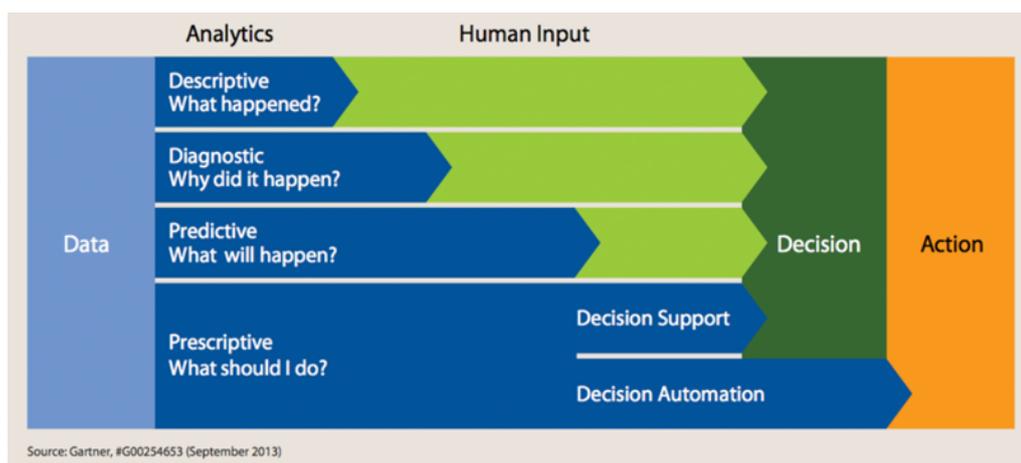


Figure 7 Machine Learning Cybernetics

References

- [1] GMC, "Global Military Communications Magazine," January 2017. [Online]. Available: <http://www.satelliteevolutiongroup.com/articles/C4ISR.pdf>. [Accessed December 2017].
- [2] K. Takeda and T. Onodera, *Artificial Intelligence: Learning Through Interactions and Big Data*, IBM Redbooks, 2013.
- [3] S. Aghabozorgi and P. Lin, "Data Scientist vs Data Engineer, What's the difference?," *Cognitive Class.ai*, 6 June 2016. [Online]. Available: <https://cognitiveclass.ai/blog/data-scientist-vs-data-engineer/>. [Accessed December 2017].
- [4] N. Callaos, "Knowlege and Cognition," March 2013. [Online]. Available: <http://www.iiis.org/Nagib%2DCallaos/Cognition%2Dand%2DKnowledge/>.
- [5] F. Maymir-Ducharme and L. Angelelli, "Cognitive Analytics: A Step Towards Tacit Knowlege?," *Journal on Systemics, Cybernetics and Informatics*, vol. JSCI Volume 12, no. Number 4, 2017.
- [6] F. Maymir-Ducharme and R. Ernst, *Optimizing Distributed and Parallel TCPED Systems, US Geospatial Intelligence Foundation (USGIF) Technical Workshop*, Denver, Colorado, 2013.
- [7] F. Maymir-Ducharme and P. Hartswick, *Advanced Cyber Defence Informatics and Cybernetics*, Rotterdam, 2017.
- [8] Sapp and Carlton, "Preparing and Architecting for Machine Learning," Gartner, 2017.